Out of Sight, Out of Mind:
Tax Salience and Mutual Fund Flows

Phorn-Arpha Phutphithak* Jing Xian†

March 27, 2026

ABSTRACT

We provide causal evidence that mutual fund investors' failure to internalize substantial tax frictions is driven by the lack of salient fund-level tax information. Exploiting Morningstar's 2005 introduction of the standardized "Tax Cost Ratio" as a shock to tax salience, we document a severe structural break in capital allocation. Utilizing a continuous difference-in-differences framework, we estimate that a one-standard-deviation increase in pre-shock tax inefficiency reduces annual net flows by 2.67 percentage points. We decompose this effect to show the market specifically penalizes uncompensated, abnormal tax drag rather than predictable, performance-driven tax liabilities. Finally, cross-sectional tests reveal these outflows are entirely concentrated among broker-distributed funds, demonstrating that new market discipline is actively enforced by professional intermediaries rather than unadvised retail investors.

JEL classification: G11, G24, G41, H24

Keywords: Mutual Fund Flow, Tax Burden, Information Salience

*Swiss Finance Institute and University of Geneva. Email:phorn-arpha.phutphithak@unige.ch

†Swiss Finance Institute and University of Geneva. Email: jing.xian@unige.ch

We thank Harald HauAndreas Neuhierland Olivier Scaillet for their comments and suggestions. We acknowledge financial support from the Swiss Finance Institute, and the Swiss National Science Foundation (SNSF Grant No. 10004770).

I. Introduction

While academic research and industry performance metrics primarily focus on pre-tax returns, this convention can obscure an economically important friction in delegated portfolio choice: the tax drag embedded in mutual fund distributions. Because mutual funds are pass-through entities, distributions of dividends and realized capital gains create immediate tax liabilities for taxable investors[1]. The economic magnitude of this friction is substantial and, for many equity mutual funds, annual tax burdens can be comparable to, or even larger than, the average expense ratio (Sialm and Zhang, 2020). Yet, despite this sizable drag on investor wealth, tax efficiency has historically received much less attention than fees, turnover, or pre-tax performance (Carhart, 1997; Wermers, 2000; Barber, Odean, and Zheng, 2005; Gil-Bazo and Ruiz-Verdú, 2009; Berk and Binsbergen, 2015). This paper studies whether making tax costs more salient changed mutual fund capital allocation, and which component of tax burden investors actually appear to penalize.

A puzzling result in the literature on taxation, returns, and fund flows is why tax-efficient mutual funds fail to attract enough capital to arbitrage away their advantage. Standard models of asset management imply that flows should compete away abnormal performance if investors internalize all relevant costs (Berk and Green, 2004). Yet tax-efficient funds have been shown to outperform less tax-efficient peers both before and after taxes, while still remaining underinvested (Sialm and Zhang, 2020). This pattern suggests that investors were not fully incorporating the tax consequences of fund choices into fund-selection decisions.

We argue that the key friction is not the complete absence of information, but the low salience of information that was already available. Tax data had long existed in prospectuses and other regulatory documents, but those disclosures were difficult to process and compare across funds. The January 2005 introduction of Morningstar's "Tax Cost Ratio" changed that environment. Morningstar condensed complex tax-distribution information

[1]Short-term capital gains are taxed at the investor's ordinary income rate, whereas long-term capital gains benefit from lower preferential statutory rates. The taxation of dividends underwent a structural shift in 2003: they were taxed as ordinary income prior to 2003, but are taxed as qualified dividends at the lower preferential rate thereafter. Additionally, a 3.8% Net Investment Income Tax applies to all such distributions starting in 2013. See Table 1 for a detailed schedule of historical federal tax rates.

into a standardized percentage that could be viewed alongside other fund characteristics on the platform most widely used by retail investors and advisors. This event did not alter the tax code or funds' underlying portfolio choices. Instead, it made tax costs easier to observe, compare, and act upon.

Our main empirical strategy features a difference-in-differences framework with a continuous treatment (Callaway, Goodman-Bacon, and Sant'Anna, 2024). First, we summarize each fund's pre-period tax inefficiency as the average abnormal tax-burden residual over 2002–2004 and ask whether post-2005 flows become more negative for funds that were more tax-inefficient before the Morningstar rollout. Second, to isolate managerial tax inefficiency from performance-chasing, we decompose pre-period tax burden into a predictable component and an abnormal component and compare their post-2005 flow effects directly. Third, we estimate difference-in-slopes specifications that test whether flows become more sensitive after 2005 to lagged raw and residual tax-burden measures. Fourth, we examine non-linearity and heterogeneity by studying the post-2005 response across quintiles of tax inefficiency and across distribution channels. Our paper therefore centers on the continuous-treatment estimates, the timing and placebo evidence, the decomposition results, the difference-in-slopes results, and the cross-sectional heterogeneity evidence.

The strongest evidence points to a meaningful post-2005 shift in how investors price tax inefficiency. In our preferred specification, funds with higher pre-period tax inefficiency experienced significantly lower flows after Morningstar's Tax Cost Ratio rollout. We further show that this result is not driven by the predictable part of tax burden. In the decomposition tests, the predictable component of tax burden is positive and significant, consistent with the fact that it is tied to performance, distributions, and other fund characteristics, whereas the abnormal component remains negative and significant. The post-2005 flow penalty is also concentrated in the upper tail of tax inefficiency and is stronger among indirect (broker) funds than among direct (no-load) funds. Finally, post-2005 flows become more negatively related to lagged raw and residual tax burden, especially when tax measures are persistent and measured over three-year windows. Taken together, the results suggest that Morningstar's standardized disclosure did not merely

repackage information: it changed which component of tax burden investors penalized in capital-allocation decisions.

Our paper contributes to three distinct strands of literature at the intersection of public economics and asset management. First, we advance the literature on limited attention and tax-aware capital allocation by demonstrating that investor sensitivity to after-tax performance is strictly conditional on information salience. Foundational work in public economics establishes that individuals systematically underreact to taxes when they are not highly salient (Chetty, Looney, and Kroft, 2009; Finkelstein, 2009). While early empirical finance literature documents that mutual fund inflows correlate with after-tax returns (Bergstresser and Poterba, 2002), we show that this market mechanism remains largely dormant when tax data are buried in complex regulatory filings. By exploiting the 2005 Morningstar disclosure shock, we provide causal evidence that market discipline emerges only when third-party intermediaries significantly reduce information processing costs, transforming publicly available but opaque data into a highly salient metric.

Second, our findings help resolve a central puzzle regarding the persistent misallocation of capital in delegated portfolio management. In standard rational models, capital flows should relentlessly chase and arbitrage away predictable performance advantages (Berk and Green, 2004). However, when markets feature severe information and search frictions, investors may systematically fail to allocate capital to the most efficient managers (Gârleanu and Pedersen, 2018). Recent empirical work highlights a glaring manifestation of this friction: highly tax-efficient funds reliably outperform their peers after taxes, yet remain systematically underinvested (Sialm and Zhang, 2020). We reconcile this disconnect by proving that the market penalty for tax drag was functionally obscured from retail investors prior to 2005. Once tax costs were made easily comparable, capital allocation aggressively penalized the abnormal component of a fund's tax burden, effectively activating the arbitrage mechanism that rational models predict.

Third, we contribute to the literature on financial intermediation and the economics of advice by showing how standardized disclosure alters market segmentation. A growing literature emphasizes that broker-distributed and direct-sold mutual funds operate as fun-

damentally distinct markets (Bergstresser, Chalmers, and Tufano, 2009; Del Guercio and Reuter, 2014). Furthermore, investors heavily rely on financial advisors or "money doctors" to reduce their cognitive load and processing costs (Gennaioli, Shleifer, and Vishny, 2015). Our cross-sectional heterogeneity results reveal that the post-2005 flow penalty for tax inefficiency is overwhelmingly concentrated in the intermediated broker channel, with retail investors in the direct channel exhibiting almost no reaction. This demonstrates that Morningstar's standardized tax metric did not uniformly awaken retail households; rather, it provided a quantitative screening heuristic that professional intermediaries and advisors utilized to enforce market discipline on behalf of their clients.

II. Data

A. Mutual Fund Sample

Our primary data source is the CRSP Survivor-Bias-Free U.S. Mutual Fund Database. We conduct our analysis at the fund level, rather than the share-class level, to capture portfolio management and tax decisions. We link share classes to portfolios using the WRDS MFLINKS table and aggregate share-class characteristics to the fund level using total net assets (TNA) lagged by one month as weights. Fund-level returns, expense ratios, and turnover are computed as TNA-weighted averages across share classes, while fund-level TNA is the sum of TNA across share classes.

Our sample period spans January 1990 to December 2024 and focuses on U.S. equity mutual funds. We exclude funds classified as fixed income, international, sector-specific, balanced, or target-date funds using CRSP objective codes and Lipper classifications. We further require funds to maintain at least 80% of their portfolio in equities for the majority of their reported history. We also exclude exchange-traded funds (ETFs), closed-end funds, and variable annuity underlying funds. To mitigate noise from micro-funds, we remove funds with less than $10 million in TNA for more than one-third of their existence. Finally, to limit incubation bias and ensure stable volatility estimates, we exclude observations where fund age is below 24 months.

We categorize the remaining U.S. equity funds into three mutually exclusive groups. *Index funds* are identified using CRSP markers for index strategies or by searching fund

names for keywords such as "Index", "S&P", "Russell", "Nasdaq", or "iShares". A fund is classified as an index fund if it meets either criterion for the majority ($> 50\%$) of its sample history. *Tax-managed funds* are identified using tax-strategy keywords in fund names (e.g., "Tax-Managed", "Tax-Efficient"). All remaining funds are classified as *actively managed* U.S. equity funds. We aggregate monthly data to an annual frequency to align with tax reporting periods, computing annual returns geometrically and averaging expense ratios and turnover. We require at least nine months of return data within a calendar year. The resulting annual panel contains roughly 50,000 fund-year observations. In the main flow analysis, we focus on actively managed funds and exclude tax-managed funds.

For descriptive exhibits, we use two closely related samples. The *broad comparison sample* retains active, index, and tax-managed funds and is used to document how tax burdens differ across product types. The main analysis focuses on an *active U.S. equity fund sample*. For exhibits tied directly to the baseline continuous-treatment design, we further restrict to a *baseline treatment sample*, defined as the subset of active U.S. equity funds with the information needed to construct the 2002–2004 residual-based tax-inefficiency measure.

We compute net fund flows, $Flow_{f,t}$, for fund f from year $t-1$ to t as the growth rate of TNA adjusted for investment returns:

$$Flow_{f,t} = \frac{TNA_{f,t} - TNA_{f,t-1}(1 + R_{f,t})}{TNA_{f,t-1}(1 + R_{f,t})}, \tag{1}$$

where $TNA_{f,t}$ denotes total net assets of fund f at the end of year t and $R_{f,t}$ is the fund's annual return. This measure captures net investor inflows or outflows, independent of portfolio performance. To mitigate the influence of extreme outliers and data errors, we winsorize flows at the 1% level in both tails in each year.

B. Distribution-Channel Classification

To study heterogeneity in investor response, we construct a pre-shock distribution-channel classification using load charges and 12b-1 fees. We classify a fund as direct (no-load) if it has no positive front- or back-end load during 2004 and its average 12b-1 fee over the same

period does not exceed 0.25%. We classify a fund as indirect (broker) if it has a positive load in 2004 or a 12b-1 fee above that threshold. This classification distinguishes funds sold through more intermediated sales channels from funds sold through lower-friction direct channels prior to Morningstar's 2005 tax-salience shock. This classification distinguishes funds sold through more intermediated sales channels from funds sold through lower-friction direct channels prior to Morningstar's 2005 tax-salience shock.

This classification follows the mutual fund distribution-channel literature that distinguishes broker-sold from direct-sold funds using load and 12b-1 fee information and treats these channels as economically distinct environments for investor search, advice, and fund demand (Bergstresser, Chalmers, and Tufano, 2009; Del Guercio and Reuter, 2014).

C. Tax Rates

Table 1 summarizes the federal tax rates we apply to mutual fund distributions over the 1990–2024 sample period. We collect the top marginal federal tax rates on dividends, short-term capital gains, and long-term capital gains from the Internal Revenue Service (IRS) Statistics of Income database. Following Sialm and Zhang (2020), we distinguish between three categories of taxable distributions. Short-term capital gains (STCG) are taxed at the ordinary income tax rate, which varies over time due to legislative changes such as the 1993 rate increases, the 2001–2003 tax cuts, and the Tax Cuts and Jobs Act of 2017. Long-term capital gains (LTCG) benefit from preferential rates, declining from 28.0% in the early 1990s to 15.0% during 2003–2012, and stabilizing at 20.0% from 2013 onward. Dividend taxation also shifts in 2003: prior to 2003, dividends are taxed as ordinary income, whereas qualified dividends are taxed at the lower LTCG rate from 2003 onward.

Following Sialm and Zhang (2020), we construct a fund-level measure of tax costs. Because mutual funds are pass-through entities, they do not pay corporate income taxes directly; instead, they distribute realized income and capital gains to shareholders, who then incur the tax liability. We decompose a fund's before-tax return, $R_{f,t}^{BT}$, into dividend yield (Y^{DIV}), short-term capital gains yield (Y^{STCG}), long-term capital gains yield

(Y^{LTCG}), and unrealized capital gains yield (Y^{UCG}):

$$R_{f,t}^{BT} = Y_{f,t}^{DIV} + Y_{f,t}^{STCG} + Y_{f,t}^{LTCG} + Y_{f,t}^{UCG}. \tag{2}$$

The after-tax return, $R_{f,t}^{AT}$, is the return net of taxes paid on taxable distributions. Taxes on unrealized gains (Y^{UCG}) are deferred until the investor liquidates their position. We define the *tax burden* (TB) as the difference between the before-tax and after-tax returns:

$$TB_{f,t} = R_{f,t}^{BT} - R_{f,t}^{AT} = \tau_t^{DIV} Y_{f,t}^{DIV} + \tau_t^{STCG} Y_{f,t}^{STCG} + \tau_t^{LTCG} Y_{f,t}^{LTCG}, \tag{3}$$

where τ_t^{DIV}, τ_t^{STCG}, and τ_t^{LTCG} are the applicable federal tax rates in year t as detailed in Table 1. This measure captures the immediate tax liability generated by a fund's taxable distributions for a long-horizon taxable shareholder. In the empirical analysis below, we use the raw tax burden to construct residual-based measures of abnormal tax inefficiency.

D. Summary Statistics

Table 2 reports summary statistics for the main samples used throughout the paper. Panel A focuses on the main descriptive sample of active U.S. equity funds. Panel B reports the baseline treatment sample used to construct the pre-period tax-inefficiency measure. Panels C through E provide benchmark comparisons for the broader active-fund universe and for the tax-managed and index-fund groups.

The summary statistics confirm that tax costs are economically meaningful in the active-fund universe. In Panel A, the mean tax burden is 1.42% per year, reducing the average before-tax return of 9.99% to an after-tax return of 8.59%. The baseline treatment sample in Panel B looks similar, with a mean tax burden of 1.44%. This similarity is reassuring: the treatment construction does not rely on a subset of funds with fundamentally different tax and return profiles than the broader active-fund sample.

Panels C through E highlight the expected ranking across product types. Actively managed funds exhibit the highest average tax burden at 1.43%, tax-managed funds the lowest at 0.44%, and index funds an intermediate burden of 1.09%. Active funds also display higher turnover and higher expense ratios, while index funds are larger and

less expensive on average. Taken together, these patterns motivate the analysis that follows: tax burden is large enough to matter for taxable investors, and there is substantial variation across funds in both tax outcomes and the characteristics that may shape those outcomes.

E. Fund-Level Tax Burden

In this section, we characterize the economic magnitude and structural properties of fund-level tax burden using the samples most relevant for the main identification. Three facts are especially important for the analysis that follows. First, average tax burdens in the baseline treatment sample of active U.S. equity funds are economically meaningful throughout the sample period. Second, there is substantial cross-sectional dispersion in tax burden even within the active-fund universe. Third, tax burden is persistent over time, implying that past tax outcomes are informative about future tax costs.

E.1. Mean Tax Burden Over Time

Figure 1 plots the average tax burden over time in the baseline treatment sample. Average tax burden is persistently positive and often economically large, typically fluctuating between roughly 1% and 2% per year, with clear spikes around episodes of strong market appreciation. This pattern highlights that tax costs are not a marginal friction for taxable investors. Instead, they are large enough to materially reduce after-tax returns in the very set of funds that underlies our main empirical design.

E.2. Cross-Sectional Dispersion of Tax Burden

Figure 2 illustrates the substantial cross-sectional variation in tax burdens within the baseline treatment sample. The spread between the upper and lower parts of the distribution remains economically meaningful throughout the sample, indicating that some funds consistently impose much larger tax liabilities on shareholders than others. This heterogeneity is central for our empirical strategy: without meaningful cross-sectional variation in tax outcomes, investor-flow responses to tax salience would be difficult to identify.

E.3. Persistence of Fund-Level Tax Burden

Figure 3 reports transition probabilities across tax-burden quintiles from year t to $t+1$. The figure shows substantial persistence, especially at the tails of the distribution. Funds in the most tax-efficient quintile frequently remain there, and funds in the least tax-efficient quintile also display pronounced persistence. Across the matrix, diagonal entries are well above the 20% benchmark implied by a random allocation. This persistence supports the central logic of the paper: tax burden is not purely transitory noise, so investors can plausibly use past tax outcomes to form expectations about future tax efficiency.

This strong persistence fundamentally validates our empirical design and ties into the broader asset pricing literature on managerial skill. As documented by Sialm and Zhang (2020), tax efficiency is not merely a product of random, one-off realization events; rather, it is a persistent, structural managerial trait. Funds in the most tax-inefficient quintile consistently utilize poor tax-management strategies, such as failing to harvest losses or systematically churning short-term winners. Because this tax drag is highly persistent, our use of the 2002–2004 pre-period accurately captures the deeply embedded "bad actors" in the cross-section. Crucially, this persistence means that past tax inefficiency is a highly informative signal for future tax drag, making it a rational metric for investors and intermediaries to penalize once it becomes salient.

E.4. Structural versus Legacy Tax Drag

A natural question when evaluating mutual fund tax burdens is whether they represent an ongoing structural friction of active management or merely a "legacy" artifact of embedded gains from previous market expansions (e.g., the 1990s bull market). Figures 4 and 5 resolve this by comparing tax burdens across fund types and cohorts.

Figure 4 demonstrates that actively managed funds impose a structurally higher tax burden than both index funds and explicitly tax-managed funds. This wedge is driven by the fundamental mechanics of active management: managers frequently trade to chase pre-tax alpha, inadvertently realizing short-term capital gains which are taxed at punishingly high ordinary income rates (as detailed in Table 1). This confirms that tax drag is an active, ongoing friction rather than a passive byproduct of holding equities.

Furthermore, Figure 5 compares the tax burdens of established incumbent funds with those of new entrants launched after 2004. If tax burdens were purely driven by legacy capital gains overhang, post-2004 entrants should exhibit near-zero tax drag. Instead, the figure shows that new entrants rapidly generate massive tax burdens that closely track the incumbents. Together, these figures confirm that severe tax inefficiency is an active, contemporary managerial choice rather than an unavoidable legacy cost, fully justifying the market's decision to penalize it.

III. Empirical Methodology

A. A Difference-in-Differences Framework with a Continuous Treatment

Our baseline design asks whether funds that were more tax-inefficient before Morningstar's 2005 disclosure shock experienced weaker flows after tax costs became easier to observe and compare. We summarize each fund's pre-period treatment intensity as the average annual abnormal tax-burden residual over 2002–2004. For each year $\tau \in \{2002, 2003, 2004\}$, we estimate the following cross-sectional first-stage regression:

$$TB_{f,\tau} = \lambda_{0,\tau} + \lambda_{1,\tau} Ret_{f,\tau-1} + \Phi_\tau Z_{f,\tau-1} + \Psi_\tau Y_{f,\tau-1} + \eta_{s,\tau} + \nu_{f,\tau}, \tag{4}$$

where $TB_{f,\tau}$ is the raw tax burden of fund f in year τ, $Ret_{f,\tau-1}$ is lagged annual return, $Z_{f,\tau-1}$ contains lagged fund characteristics, and $\eta_{s,\tau}$ denotes style fixed effects estimated separately by year. The first-stage controls include lagged log fund size, lagged log fund age, lagged expense ratio, lagged turnover, lagged volatility, and lagged distribution yields for dividends, short-term capital gains, and long-term capital gains. The fitted value from Equation (4) captures the predictable part of tax burden, while the residual $\nu_{f,\tau}$ captures abnormal tax burden conditional on lagged performance, fund characteristics, lagged distribution yields, and broad style.

We then average the annual residuals over 2002–2004 to construct a time-invariant treatment-intensity measure:

$$TI_f = \frac{1}{3} \sum_{\tau=2002}^{2004} \hat{\nu}_{f,\tau}. \tag{5}$$

This measure is our main proxy for persistent pre-period tax inefficiency. We relate it to post-2005 fund flows using the following reduced-form specification:

$$Flow_{f,t} = \alpha_f + \gamma_{s\times t} + \beta(TI_f \times Post_t) + \theta X_{f,t-1} + \epsilon_{f,t}, \tag{6}$$

where $Flow_{f,t}$ denotes the net flow to fund f in year t, $Post_t$ equals one for 2005 and later, α_f are fund fixed effects, and $\gamma_{s\times t}$ are style-by-year fixed effects. The vector $X_{f,t-1}$ includes lagged return, lagged return squared, lagged log fund size, lagged log fund age, lagged expense ratio, lagged turnover, lagged volatility, and lagged flow. Standard errors are clustered at the fund level. We interpret β as a summary measure of differential post-2005 reallocation across funds with different levels of pre-period abnormal tax inefficiency.

B. Timing and Placebo Design

To assess whether the estimated break is specific to the 2005 Morningstar salience shock, we re-estimate the main specification using alternative event years and matched rolling three-year pre-period windows. For a candidate break year T, we construct the treatment intensity from the three years ending in $T-1$ and estimate the same fund fixed effects plus style×year fixed effects specification on a common post indicator defined relative to that candidate break. The placebo design therefore compares four windows: Post-2002, Post-2003, Post-2004, and Post-2005.

This timing exercise serves two purposes. First, the Post-2002 specification captures the information-availability benchmark associated with the SEC disclosure regime. Second, the Post-2003 and Post-2004 specifications are placebo breaks with no corresponding Morningstar salience change. A strong effect only in the Post-2005 specification supports the interpretation that the main result reflects salience rather than a generic trend in tax-related information.

We do not run an additional placebo for years after 2005 because those later years are already inside the treated regime. Once the Morningstar redesign is in place, a later break date would compare treated years with other treated years rather than with a clean pre-treatment benchmark, so it would no longer isolate the original disclosure shock.

C. *Raw, Predictable, and Abnormal Tax Burden*

A central issue in this setting is that raw tax burden mixes economically different components. Part of tax burden is predictable from lagged performance, lagged fund characteristics, and lagged distribution yields, while another part reflects tax inefficiency beyond those observables. In plain terms, the predictable component is the part investors might expect given recent returns and fund characteristics, whereas the abnormal component is the residual tax drag beyond those observables. To separate these channels, we use the first-stage fitted values from Equation (4) to construct the predictable component of tax burden and compare it directly with the abnormal component in the second-stage regressions. This comparison is informative because the predictable component can capture performance-related tax burden that is likely to be correlated with return-chasing, whereas the abnormal component is closer to tax inefficiency in the sense relevant for the post-2005 reallocation tests.

D. *Mechanism: Post-2005 Changes in Flow Sensitivity to Tax Measures*

To study mechanism, we estimate a difference-in-slopes specification that interacts post-2005 with lagged tax measures:

$$Flow_{f,t} = \alpha_f + \gamma_{s\times t} + \beta_1 M_{f,t-1} + \beta_2 \left(M_{f,t-1} \times Post_{2005,t}\right) + \theta X_{f,t-1} + \epsilon_{f,t}. \tag{7}$$

Here, $M_{f,t-1}$ is defined alternatively as the one-year lagged raw tax burden, the three-year moving average of lagged raw tax burden, the one-year lagged residual tax burden, or the three-year moving average of lagged residual tax burden. The coefficient β_1 captures the pre-2005 flow sensitivity to lagged tax measures, while β_2 captures the incremental change in that sensitivity after 2005. We interpret a negative β_2 as evidence that, after Morningstar introduced the Tax Cost Ratio, investor flows became more negatively related to recent tax burden or recent abnormal tax burden.

E. *Distribution Channel Heterogeneity*

We study whether the post-2005 reallocation is stronger in funds distributed through financial intermediaries. We classify funds as direct (no-load) or indirect (broker) using sales-load and 12b-1 information measured in 2004. Direct funds have no positive sales load in 2004 and 12b-1 fees no greater than 25 basis points, while indirect funds have a positive sales load in 2004 or a 12b-1 fee above that threshold. This distinction is economically meaningful because broker-sold and direct-sold funds differ in how investors receive advice and recommendations, and prior work shows that intermediary incentives can shape which funds attract flows in broker channels (Bergstresser, Chalmers, and Tufano, 2009; Christoffersen, Evans, and Musto, 2013).

We first estimate the baseline specification separately within the direct and indirect subsamples. This split-sample approach provides a transparent comparison of whether the post-2005 flow penalty is concentrated in one distribution channel. We then estimate a pooled triple-interaction specification:

$$Flow_{f,t} = \alpha_f + \gamma_{s\times t} + \beta_1(TI_f \times Post_t) + \beta_2(TI_f \times Post_t \times Indirect_f) + \theta X_{f,t-1} + \epsilon_{f,t}, \quad (8)$$

where $Indirect_f$ equals one for broker-distributed funds. Because $Indirect_f$ is time-invariant, its main effect is absorbed by the fund fixed effects. In this specification, β_1 captures the post-2005 penalty among direct funds, and β_2 captures the incremental penalty for indirect funds relative to direct funds.

F. *Dose-Response and Non-Linearity*

To study whether the post-2005 flow response is concentrated among the most tax-inefficient funds, we estimate a flexible dose-response specification. We sort funds into quintiles based on the pre-period abnormal tax-burden measure TI_f and interact each quintile indicator with the post-2005 dummy, using the lowest-tax-inefficiency quintile as the omitted group:

$$Flow_{f,t} = \alpha_f + \gamma_{s\times t} + \sum_{q=2}^{5} \beta_q \left(\mathbb{1}\{TI_f \in Q_q\} \times Post_t\right) + \theta X_{f,t-1} + \epsilon_{f,t}. \quad (9)$$

This specification relaxes the linearity embedded in Equation (6) and allows us to examine whether the post-2005 response is concentrated in the upper tail of tax inefficiency rather than spread uniformly across the entire distribution.

G. Robustness Design

We implement two main robustness exercises. First, we re-estimate the preferred difference-in-differences specification on the narrower 2000–2012 window. This subperiod keeps the analysis centered on the years closest to the 2005 disclosure change while preserving the same preferred fixed effects and lagged controls used in the main tables. We report these estimates for the pooled sample and for the direct and indirect subsamples.

Second, we examine whether the channel results depend on the exact broker classification rule. We therefore re-estimate the broker-channel specification under alternative definitions based on front-end loads, any load in 2004, 12b-1 thresholds, high-load cutoffs, and an average 2002–2004 rule. If the indirect-fund result remains strong across these definitions, the channel interpretation is unlikely to be an artifact of one specific classification threshold.

IV. Empirical findings

A. The Causal Effect of Tax Salience on Mutual Fund Flows

Table 3 reports the baseline continuous difference-in-differences estimates of annual net fund flows. The coefficient of interest is the interaction between pre-period tax inefficiency and the post-2005 indicator, which captures the differential flow response to abnormal tax burdens following the Morningstar disclosure shock. By defining the treatment variable as the average abnormal tax-burden residual over the 2002–2004 pre-window, the baseline specification separates tax inefficiency from tax liabilities that are mechanically related to past performance.

Across the three fixed-effects specifications, the coefficient on pre-period tax inefficiency $\times$ Post-2005 remains consistently negative and increases in magnitude as the fixed-effects structure becomes richer. In our preferred specification in Column (3), which incorporates both fund fixed effects and style-by-year fixed effects, the coefficient reaches

−5.748 and is statistically significant at the 5% level. This increase in magnitude is economically revealing: because different investment styles, such as growth versus value, experience distinct performance cycles that mechanically drive both capital gains realizations and investor flows, static style controls in Column (2) are insufficient. By fully absorbing these time-varying, style-specific market shocks, Column (3) ensures the penalty reflects a strict within-style reallocation away from tax-inefficient funds after tax costs became salient.

To contextualize the economic magnitude of the 2.67 percentage point reduction in annual net flows, it is instructive to evaluate the implied aggregate capital reallocation. Within our baseline sample, the average total net assets (TNA) of an actively managed equity fund is approximately $2.1 billion. A 2.67 percentage point flow penalty therefore translates to an average annual outflow of roughly $56 million per fund for a one-standard-deviation increase in tax inefficiency. Extrapolating this penalty across the upper quintiles of the tax-inefficiency distribution suggests that the Morningstar disclosure shock catalyzed the reallocation of tens of billions of dollars annually. This aggregate magnitude underscores that information processing costs previously masked a friction of first-order importance to the asset management industry, and that mitigating these costs substantially altered the equilibrium distribution of capital.

B. Timing Validation: Why the Break Is 2005

Table 4 evaluates whether the break in the baseline specification is specific to 2005. The pooled timing design compares the actual 2005 salience break with an information-availability benchmark in 2002 and placebo breaks in 2003 and 2004, each paired with a matched rolling three-year treatment window.

The estimates show a clear timing pattern. The Post-2002 and Post-2003 coefficients are small and statistically insignificant. The Post-2004 estimate is positive and only marginally significant, rather than resembling the negative effect observed in 2005. Column (4) then reports the preferred full-sample 2005 benchmark using the same all-funds specification as Table 3. Its coefficient remains negative and statistically significant. Taken together, these results support the interpretation that the main break is tied to the 2005 Morningstar salience change rather than to a general trend in tax disclosure or

to a gradual pre-existing shift in investor demand.

The marginally positive coefficient in the Post-2004 placebo further confirms the absence of any negative pre-trend; investors were absolutely not penalizing tax inefficiency before the Morningstar tax information disclosure. Taken together, these results support the interpretation that the structural break is uniquely tied to the 2005 salience change rather than a gradual shift in investor demand.

C. Decomposing the Flow Penalty: Raw, Predictable, and Abnormal Tax Burdens

Table 5 isolates the precise economic channel driving the post-2005 capital reallocation by decomposing the pre-period tax burden into its raw, predictable, and abnormal components. A fundamental econometric challenge in evaluating tax salience is omitted variable bias: raw tax burdens mechanically mix two opposing forces. Because strong past returns inevitably generate high capital gains realizations, the predictable component of the tax burden acts as a proxy for historical outperformance. Given that retail investors are notorious performance-chasers, failing to decompose the tax burden risks confounding the negative penalty for tax inefficiency with the positive flow response to strong returns.

The empirical results in Table 5 perfectly align with this economic intuition. On its own, the raw tax burden yields an economically small and statistically insignificant coefficient, as it conflates the reward for returns with the penalty for taxes. Once decomposed, these forces cleanly separate. The predictable component enters with a strongly positive and highly significant coefficient , successfully capturing the well-documented performance-chasing behavior of retail investors.

Conversely, the abnormal component, which captures the pure, uncompensated deadweight loss of poor tax management orthogonal to returns and strategy, yields a strictly negative and highly significant coefficient. This decomposition provides critical insight into the nature of the Morningstar shock. Before 2005, opacity prevented the market from disentangling good performance from bad tax management. By standardizing this data, the salience shock allowed investors to selectively and sophisticatedly punish uncompensated abnormal tax drag while continuing to reward high pre-tax returns.

D. *Difference-in-Slopes: Post-2005 Changes in Flow-Tax Sensitivity*

Table 6 complements the baseline difference-in-differences design with a difference-in-slopes specification that uses time-varying lagged tax measures. Each column reports the pre-2005 slope on the lagged tax measure and the incremental change in that slope after 2005.

Across all four specifications, the pre-2005 slope is positive, which is consistent with investors rewarding funds whose higher tax burdens partly reflect strong past performance. The post-2005 interaction is negative in every case and becomes especially large for the three-year measures. The strongest result appears for the three-year abnormal tax burden measure, where the post-2005 slope change is -5.428. This indicates that the post-2005 reallocation is most pronounced for persistent tax inefficiency rather than for one-off tax events.

E. *Heterogeneity by Distribution Channel: Direct and Indirect Funds*

To deepen our understanding of the economic mechanism driving the post-2005 flow penalty, Tables 7 and 8 examine whether the sensitivity to tax salience differs across fund distribution channels. A central question in the limited attention literature is *who* actually processes and reacts to newly salient disclosures. To test this, we partition our sample into direct-sold (no-load) funds, typically purchased by do-it-yourself (DIY) retail investors, and indirect funds, which levy distribution fees and are sold through third-party financial advisors and broker-dealers.

The split-sample estimates in Table 7 and the pooled triple-interaction in Table 8 reveal striking cross-sectional heterogeneity. The post-2005 flow penalty is entirely concentrated in the indirect, broker-distributed channel, yielding a massive and highly significant additional outflow of 12.060 percentage points per unit of treatment intensity. In contrast, the coefficient for direct-sold funds is economically small and statistically indistinguishable from zero.

This divergence aligns seamlessly with the institutional realities of the mutual fund industry. The direct-sold channel is predominantly populated by unadvised retail investors who face substantial search and information processing costs. Even after Morningstar standardized the Tax Cost Ratio, a DIY investor still had to actively search for a fund

and manually compare secondary metrics. Overcoming this cognitive friction to optimize over after-tax characteristics often exceeds the average retail investor's attention threshold, rendering the standardized disclosure alone insufficient to alter capital allocation.

Conversely, the indirect channel is characterized by professional financial intermediation. Advisors systematically utilize professional analytics platforms, notably Morningstar Direct, to construct client portfolios and justify their advisory fees. Within this intermediated framework, the introduction of the new tax metric provided a quantifiable, standardized heuristic that advisors could efficiently deploy at scale to screen out underperforming managers on an after-tax basis.

Ultimately, the market discipline documented in this paper was not driven uniformly by retail households waking up to taxes, but was actively enforced by professional intermediaries leveraging the new disclosure to optimize client portfolios. This demonstrates that salience shocks are most effective when they intersect with agents possessing the institutional capacity to process them, consistent with prior evidence that broker-sold fund demand uniquely reflects intermediary incentives rather than pure retail preferences (Bergstresser, Chalmers, and Tufano, 2009; Christoffersen, Evans, and Musto, 2013; Del Guercio and Reuter, 2014).

F. Non-Linearity in the Flow Penalty: Thresholds and Extreme Salience

Our baseline continuous-treatment specification assumes a linear relationship between tax inefficiency and subsequent capital outflows. However, theories of limited attention and bounded rationality suggest that investor responses to new disclosures are often highly non-linear due to attention thresholds and switching costs. Figure 6 explores this by plotting the dose-response estimates across quintiles of pre-period abnormal tax inefficiency, using the most tax-efficient funds (Quintile 1) as the baseline.

The empirical pattern reveals a stark non-linearity. The post-2005 flow penalty is small and statistically insignificant for funds in the second quintile, but becomes aggressively negative in the upper tail of the distribution, entirely concentrated in Quintiles 3 through 5.

This shape provides deep economic insight into how salience practically alters capital allocation. Intermediaries and investors do not ruthlessly optimize over every single basis

point of tax drag. Moving from a highly efficient fund to a moderately efficient one (e.g., Quintile 1 to Quintile 2) does not trigger capital flight because the marginal tax drag fails to breach the investor's attention threshold and cannot justify the tangible switching costs (such as capital gains realized upon selling the old fund). However, severe tax inefficiency acts as a glaring red flag. Once Morningstar's salience shock standardized the data, intermediaries utilized the Tax Cost Ratio as a negative screening heuristic, systematically purging their clients' portfolios of the most egregious, wealth-destroying outliers in the top quintiles.

G. Alternative Explanations and Confounding Macroeconomic Events

While the baseline estimates attribute the structural break in fund flows to the 2005 Morningstar disclosure shock, an alternative hypothesis is that the reallocation reflects a delayed investor response to the Jobs and Growth Tax Relief Reconciliation Act (JGTRRA) of 2003 as shown in Table 1, which significantly reduced statutory tax rates on dividends and long-term capital gains. However, the empirical patterns documented in our analysis are inconsistent with a purely legislative explanation. If the flow penalty were driven primarily by the 2003 tax cuts, the reallocation should logically begin in 2003 or 2004. Our timing and placebo tests (Table—4) explicitly reject this hypothesis, demonstrating that the flow penalty remains statistically indistinguishable from zero in the Post-2003 and Post-2004 placebo specifications, only materializing precisely in 2005.

Furthermore, if statutory rate reductions were the primary driver of investor behavior, the penalty for tax inefficiency should theoretically attenuate rather than intensify, as the 2003 legislation lowered the maximum tax burden on capital distributions. The fact that the market penalty strictly increases after 2005, despite the lower statutory tax environment, supports the hypothesis that the binding constraint on investor behavior was information salience rather than the absolute level of the tax schedule.

H. Robustness

The main results remain stable in two additional exercises. First, Table 9 re-estimates the preferred specification on the narrower 2000–2012 window for the pooled sample and for the direct and indirect subsamples. The pooled and indirect-fund effects remain negative

and become larger in magnitude in the narrower window, while the direct-fund effect remains statistically insignificant. This shows that the main findings are not driven by distant sample years outside the period closest to the Morningstar redesign.

Second, Table 10 shows that the indirect-fund result is robust to alternative ways of defining broker-distributed funds. The estimated post-2005 penalty remains negative and statistically significant across all six classification rules. This consistency reduces the concern that the channel results depend on one arbitrary threshold for loads or 12b-1 fees.

V. Conclusion

This paper provides evidence that tax salience matters for how investors allocate capital across mutual funds. We show that funds with higher pre-period tax inefficiency experience weaker post-2005 flows after Morningstar's 2005 Tax Cost Ratio made tax costs easier to observe and compare. We further show that this post-2005 penalty is driven by the abnormal component of tax burden rather than by the predictable component linked to returns, distributions, and other fund characteristics.

Our interpretation is not that investors directly observed the residualized tax-inefficiency measure used in the baseline DiD. Rather, the residual serves as a proxy for persistent abnormal tax inefficiency, while the decomposition and slope-change results show that investors differentiated between predictable tax burden and abnormal tax burden after 2005. The difference-in-slopes evidence indicates that post-2005 flows became more negatively related to lagged tax measures, especially persistent multi-year raw and residual tax burden. The heterogeneity and dose-response evidence further suggest that the reallocation away from tax-inefficient funds was strongest in broker-distributed funds and among funds in the upper tail of pre-period tax inefficiency.

These findings have broader implications for disclosure policy and market design. Public availability of information is insufficient if relevant data remain buried in technical filings with high processing costs. In our setting, a third-party platform transformed opaque tax data into a standardized, salient metric. However, our channel heterogeneity results emphasize that the resulting market discipline was not driven uniformly by retail

investors, but rather enforced by financial intermediaries who leveraged the new metric to screen funds. Ultimately, our results suggest that for disclosure to effectively alter capital allocation, it must not only reduce processing costs, but do so in a format readily adoptable by the intermediaries who guide investor choice.

REFERENCES

Barber, Brad M, Terrance Odean, and Lu Zheng, 2005, Out of sight, out of mind: the effects of expenses on mutual fund flows, *The Journal of Business* 78, 2095–2120.

Bergstresser, Daniel, John M. R. Chalmers, and Peter Tufano, 2009, Assessing the costs and benefits of brokers in the mutual fund industry, *Review of Financial Studies* 22, 4129–4156.

Bergstresser, Daniel and James Poterba, 2002, Do after-tax returns affect mutual fund inflows?, *Journal of Financial Economics* 63, 381–414.

Berk, Jonathan and Richard Green, 2004, Mutual fund flows and performance in rational markets, *Journal of Finance* 59, 1921–1955.

Berk, Jonathan B. and Jules H. van Binsbergen, 2015, Measuring skill in the mutual fund industry, *Journal of Financial Economics* 118, 1–20.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro HC Sant'Anna (2024). *Difference-in-differences with a continuous treatment*. Tech. rep. National Bureau of Economic Research.

Carhart, Mark M., 1997, On persistence in mutual fund performance, *The Journal of Finance* 52, 57–82.

Chetty, Raj, Adam Looney, and Kory Kroft, 2009, Salience and taxation: theory and evidence, *The American Economic Review* 99, 1145–1177.

Christoffersen, Susan E. K., Richard Evans, and David K. Musto, 2013, What do consumers' fund flows maximize? evidence from their brokers' incentives, *The Journal of Finance* 68, 201–235.

Del Guercio, Diane and Jonathan Reuter, 2014, Mutual fund performance and the incentive to generate alpha, *The Journal of Finance* 69, 1673–1704.

Finkelstein, Amy, 2009, E-ZTax: tax salience and tax rates, *The Quarterly Journal of Economics* 124, 969–1010.

Gârleanu, Nicolae and Lasse Heje Pedersen, 2018, Efficiently inefficient markets for assets and asset management, *American Economic Review* 108, 1663–1712.

Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny, 2015, Money doctors, *The Journal of Finance* 70, 91–114.

Gil-Bazo, Javier and Pablo Ruiz-Verdú, 2009, The relation between price and performance in the mutual fund industry, *The Journal of Finance* 64, 2153–2183.

Sialm, Clemens and Hanjiang Zhang, 2020, Tax-efficient asset management: evidence from equity mutual funds, *The Journal of Finance* 75, 735–777.

Wermers, Russ, 2000, Mutual fund performance: an empirical decomposition into stock-picking talent, style, transaction costs, and expenses, *The Journal of Finance* 55, 1655–1695.

I. Figures

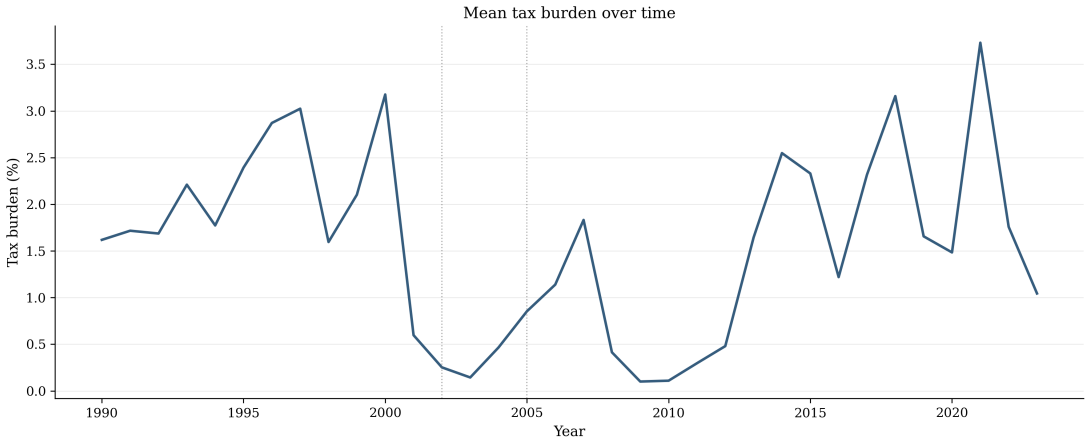


Figure 1: Mean tax burden over time

This figure plots the annual mean tax burden in the baseline active-fund sample used to construct the pre-period treatment variable. Tax burden is measured in percentage points and is defined as the difference between before-tax and after-tax fund returns.

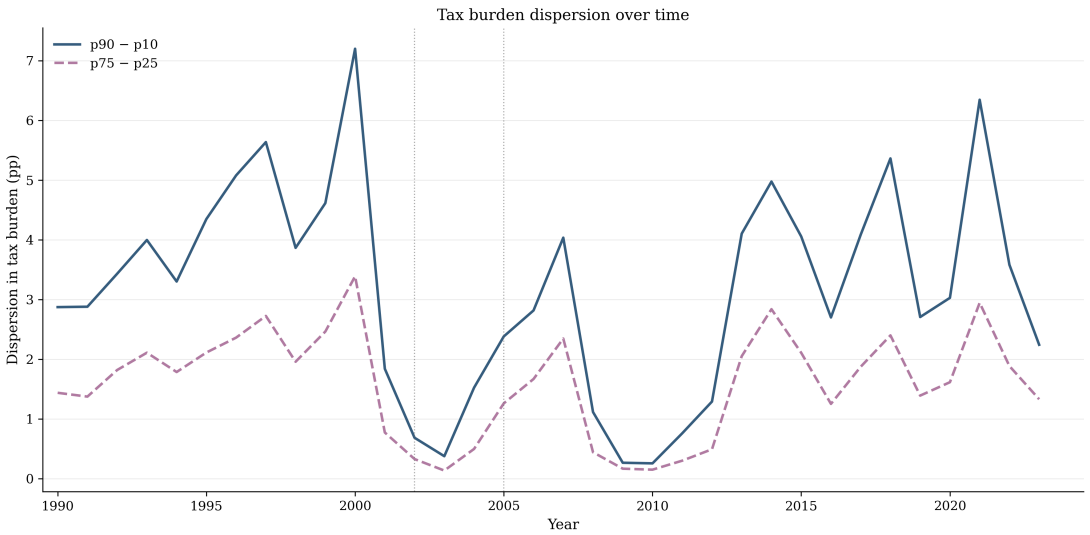


Figure 2: Cross-sectional dispersion of tax burden over time

This figure reports two spread measures for fund-level tax burden in the baseline active-fund sample. The solid line reports the difference between the 90th and 10th percentiles of the yearly cross-sectional distribution, and the dashed line reports the interquartile range. Both series indicate substantial cross-sectional heterogeneity in tax burden through time.

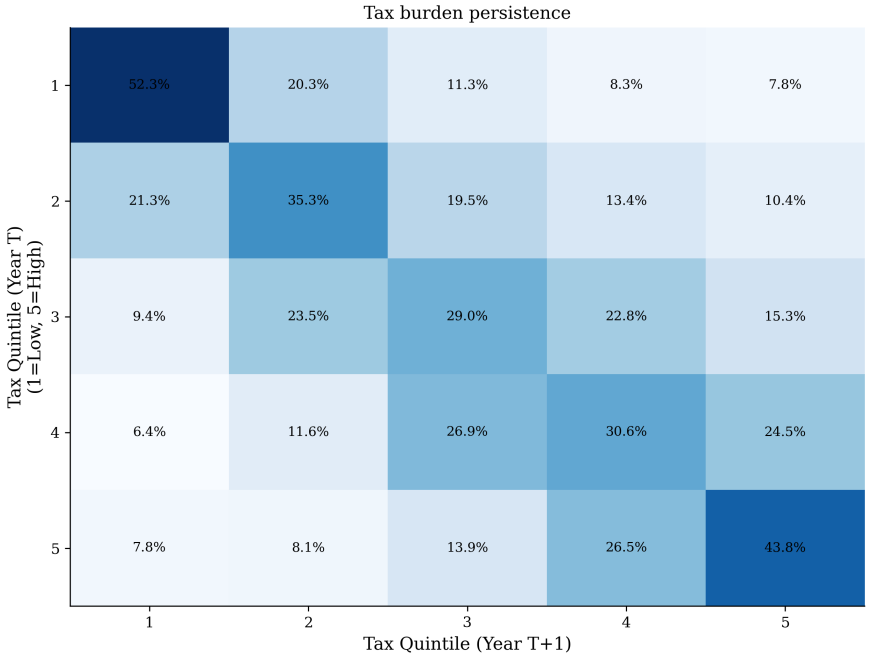


Figure 3: Persistence of fund-level tax burden

This figure reports one-year transition probabilities across tax-burden quintiles in the baseline active-fund sample. Each cell gives the share of funds moving from quintile q_t in year t to quintile q_{t+1} in year $t+1$.

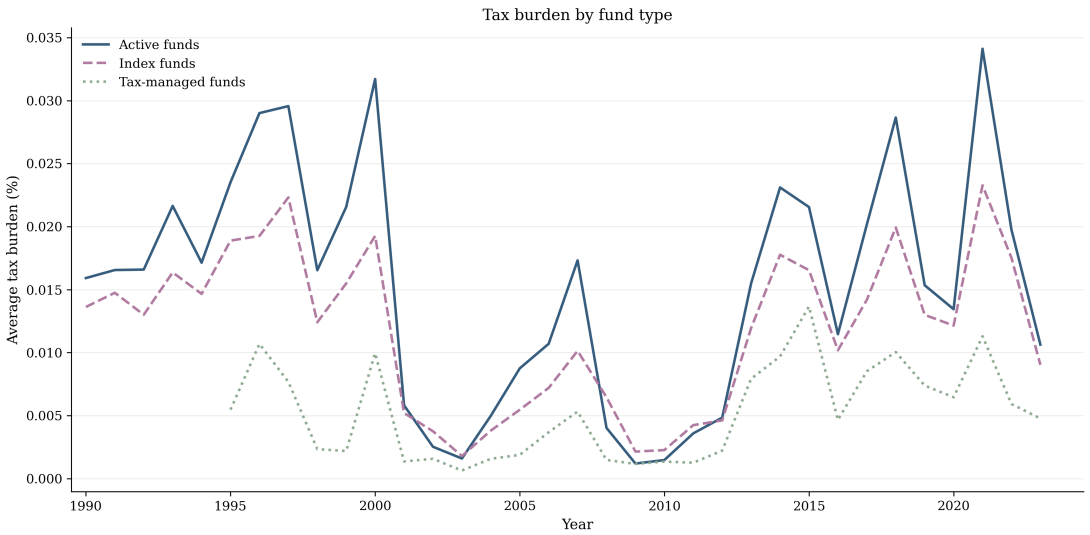


Figure 4: Tax burden by fund type

This figure compares the average tax burden of active, index, and tax-managed funds in the broad comparison sample.

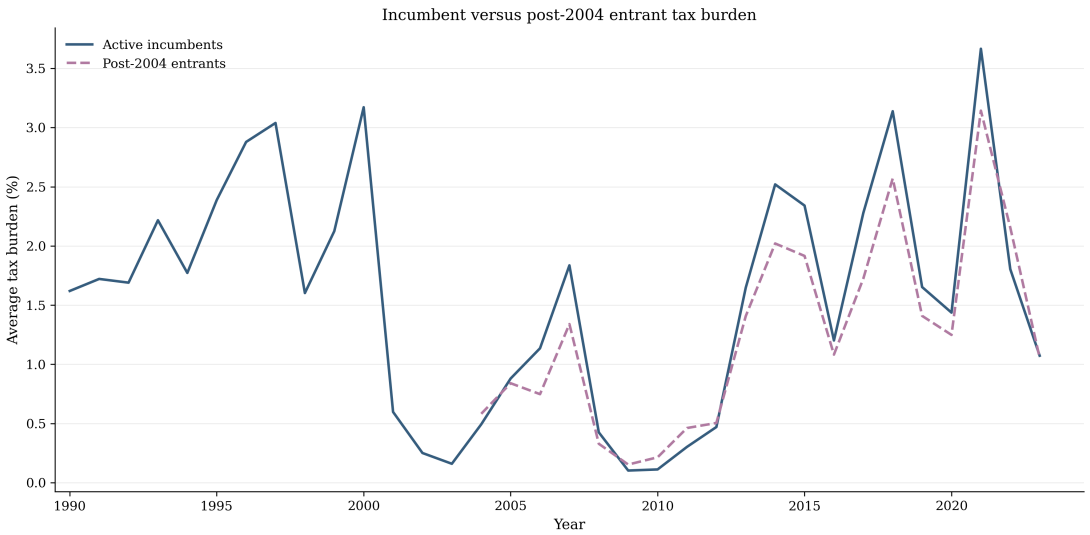


Figure 5: Incumbent versus post-2004 entrant tax burden

This figure compares the average tax burden of established active funds with that of active funds entering after 2004.

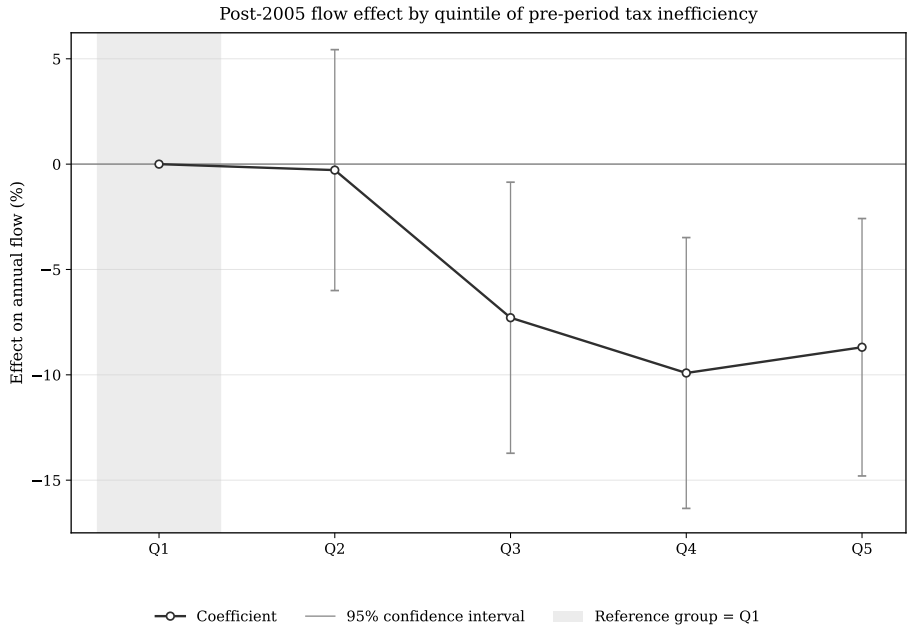


Figure 6: Dose response of the post-2005 flow effect

This figure reports the differential post-2005 flow effect across quintiles of pre-period tax inefficiency. Quintile 1 is the omitted group. The coefficients are estimated from a specification with fund fixed effects, style×year fixed effects, and the same lagged controls as in Table 3. Vertical bars show 95% confidence intervals, and standard errors are clustered at the fund level.

II. Tables

Table 1: **Tax Rates on Mutual Fund Distributions**
This table reports the tax rates applied to mutual fund distribution components: short-term capital gains (STCG), long-term capital gains (LTCG), and dividends, following Sialm and Zhang (2020). STCG distributions are taxed at ordinary income rates. LTCG distributions are taxed at preferential rates (28% prior to 1997; 20% from 1997–2002; 15% from 2003–2012; and 20% thereafter). Dividends are taxed as ordinary income prior to 2003 and at the qualified (LTCG) rate from 2003 onward. For years starting in 2013, the 3.8% Net Investment Income Tax is added.

Period	**Ordinary Income (STCG)**	**Long Term Cap Gains (LTCG)**	**Qualified Dividends**
1990	28.0%	28.0%	28.0%
1991–1992	31.0%	28.0%	31.0%
1993–1996	39.6%	28.0%	39.6%
1997–2000	39.6%	20.0%	39.6%
2001–2002	39.1% / 38.6%	20.0%	39.1% / 38.6%
2003–2012	35.0%	15.0%	15.0%
2013–2017	39.6%	20.0%	20.0%
2018–2024	37.0%	20.0%	20.0%

Table 2: **Summary Statistics**
This table reports summary statistics for the main descriptive samples and broad benchmark fund groups. *Tax Burden* is the difference between before-tax and after-tax returns. *Before-Tax Return* is the geometric annual fund return. *After-Tax Return* adjusts for taxes on dividends and capital gains distributions. *Dividend*, *short-term capital gains (STCG)*, and *long-term capital gains (LTCG)* are the component yields of the fund's distributions. *TNA* is total net assets in millions. *Expense Ratio* and *Turnover* are annual fund-level measures. *Age* is years since inception. *Average Flow* is the annual percentage growth in net assets due to new investments. *Volatility* is the annualized standard deviation of monthly returns. Panel A reports the main descriptive sample of active U.S. equity funds. Panel B reports the baseline treatment sample used to construct pre-period tax inefficiency. Panels C–E are reported on the next page.

Variable	**Count**	**Mean**	**Std. Dev.**	**Q1**	**Median**	**Q3**
Panel A: Main descriptive sample – Active U.S. equity funds						
Tax Burden (%)	31,858	1.42	1.82	0.09	0.74	2.11
Before-Tax Return (%)	32,086	9.99	19.80	-0.72	11.66	23.28
After-Tax Return (%)	31,970	8.59	19.72	-2.35	10.28	21.69
Dividend Yield (%)	31,970	0.69	1.19	0.00	0.30	1.05
STCG Yield (%)	31,970	0.95	3.11	0.00	0.00	0.46
LTCG Yield (%)	31,970	4.38	7.21	0.00	1.61	6.56
Expense Ratio (%)	32,005	1.17	0.38	0.92	1.13	1.39
Turnover (%)	32,041	75.99	70.81	31.33	57.17	96.67
TNA ($ Millions)	32,086	1,986.10	7,135.66	102.58	354.47	1,265.24
Age (Years)	32,086	17.22	14.13	7.44	13.50	22.42
Average Flow (%)	31,974	8.46	57.16	-15.56	-5.52	11.47
Volatility (%)	32,086	16.07	7.53	10.71	14.58	20.40
Panel B: Baseline treatment sample (Active U.S. equity)						
Tax Burden (%)	29,109	1.44	1.84	0.09	0.76	2.15
Before-Tax Return (%)	29,220	9.98	19.88	-0.97	11.62	23.42
After-Tax Return (%)	29,220	8.54	19.79	-2.62	10.18	21.75
Dividend Yield (%)	29,220	0.69	1.21	0.00	0.30	1.05
STCG Yield (%)	29,220	0.95	3.16	0.00	0.00	0.44
LTCG Yield (%)	29,220	4.40	7.15	0.00	1.62	6.62
Expense Ratio (%)	29,157	1.17	0.39	0.91	1.13	1.39
Turnover (%)	29,185	76.23	71.50	31.33	57.42	96.50
TNA ($ Millions)	29,220	2,105.83	7,435.63	110.25	379.00	1,355.27
Age (Years)	29,220	17.92	14.49	8.00	14.08	23.35
Average Flow (%)	29,109	7.76	55.55	-15.41	-5.57	11.00
Volatility (%)	29,220	16.14	7.59	10.77	14.68	20.44

Table 2 (continued). Panels C–E report broad benchmark comparisons for Active, Tax-Managed, and Index funds.

Variable	Count	Mean	Std. Dev.	Q1	Median	Q3
Panel C: All Active U.S. equity funds						
Tax Burden (%)	44,964	1.43	1.83	0.11	0.75	2.11
Before-Tax Return (%)	45,802	10.26	19.43	-0.72	11.98	23.54
After-Tax Return (%)	45,364	8.82	19.39	-2.31	10.58	21.94
Dividend Yield (%)	45,364	0.74	1.34	0.00	0.37	1.12
STCG Yield (%)	45,364	0.92	3.08	0.00	0.00	0.44
LTCG Yield (%)	45,364	4.48	7.41	0.00	1.71	6.63
Expense Ratio (%)	42,323	1.15	0.39	0.90	1.10	1.37
Turnover (%)	42,460	75.24	72.43	31.00	56.08	94.83
TNA ($ Millions)	45,802	1,666.28	6,172.81	90.80	302.16	1,064.88
Age (Years)	45,802	15.24	12.97	6.33	11.75	20.00
Average Flow (%)	45,389	8.35	57.48	-15.62	-5.49	11.52
Volatility (%)	45,802	16.16	7.51	10.88	14.70	20.46
Panel D: Tax-managed funds						
Tax Burden (%)	604	0.44	0.92	0.02	0.17	0.42
Before-Tax Return (%)	624	7.54	18.85	-2.78	10.50	20.73
After-Tax Return (%)	604	7.39	18.72	-2.45	10.14	20.48
Dividend Yield (%)	604	0.61	0.62	0.00	0.45	1.07
STCG Yield (%)	604	0.09	0.82	0.00	0.00	0.00
LTCG Yield (%)	604	1.26	3.42	0.00	0.00	0.37
Expense Ratio (%)	553	1.18	0.48	0.83	1.11	1.55
Turnover (%)	623	62.74	65.60	21.75	44.36	79.17
TNA ($ Millions)	624	1,170.89	2,468.31	82.35	263.80	919.32
Age (Years)	624	10.99	6.96	5.17	9.17	16.02
Average Flow (%)	624	7.65	50.21	-12.18	-2.07	10.29
Volatility (%)	624	15.26	6.77	10.49	13.70	20.25
Panel E: Index funds						
Tax Burden (%)	6,020	1.09	1.43	0.22	0.54	1.47
Before-Tax Return (%)	6,143	10.54	18.50	0.25	13.31	24.49
After-Tax Return (%)	6,092	9.40	18.57	-1.12	12.19	23.16
Dividend Yield (%)	6,092	1.15	1.25	0.41	1.09	1.60
STCG Yield (%)	6,092	0.56	2.23	0.00	0.00	0.24
LTCG Yield (%)	6,092	3.03	5.82	0.00	0.47	4.19
Expense Ratio (%)	4,304	0.87	0.47	0.46	0.74	1.20
Turnover (%)	5,133	57.17	89.78	14.00	25.00	61.83
TNA ($ Millions)	6,143	7,459.53	57,107.18	142.80	541.70	1,963.40
Age (Years)	6,143	13.64	9.71	6.25	11.67	18.92
Average Flow (%)	6,070	9.14	50.19	-11.77	-1.32	13.69
Volatility (%)	6,143	15.78	7.38	10.67	14.33	20.54

Table 3: **Investor Flows and the 2005 Tax-Salience Shock**
This table reports baseline continuous difference-in-differences estimates of annual net fund flows. The treatment variable is pre-period tax inefficiency, defined as the average abnormal tax-burden residual over 2002–2004 and expressed in percentage points. The three columns report progressively richer fixed-effects specifications from the pooled full sample of active U.S. equity funds. All columns include the same lagged controls: lagged return, lagged return squared, lagged log fund size, lagged log fund age, lagged expense ratio, lagged turnover, lagged volatility, and lagged flow. Standard errors are clustered at the fund level. t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)
Pre-period tax inefficiency × Post-2005	-5.212*	-5.217*	-5.748**
	(-1.724)	(-1.711)	(-1.980)
Lagged return (%)	1.021***	1.021***	1.452***
	(20.309)	(20.146)	(19.103)
Lagged return squared	-0.004***	-0.004***	-0.007***
	(-4.765)	(-4.744)	(-7.180)
Lagged log fund size	-16.770***	-16.707***	-16.827***
	(-21.328)	(-21.397)	(-21.760)
Lagged log fund age	-5.420***	-5.309***	-5.339***
	(-4.457)	(-4.400)	(-4.502)
Lagged expense ratio (%)	-8.844***	-8.774***	-9.194***
	(-3.182)	(-3.149)	(-3.380)
Lagged turnover	1.011	1.046	1.174
	(1.036)	(1.069)	(1.206)
Lagged volatility (%)	0.480***	0.489***	0.432***
	(4.318)	(4.359)	(2.875)
Lagged flow (%)	0.207***	0.206***	0.201***
	(13.597)	(13.537)	(13.602)
Fund FE	Yes	Yes	Yes
Style FE	No	Yes	No
Year FE	Yes	Yes	No
Style×Year FE	No	No	Yes
Observations	27,663	27,663	27,663
Unique funds	1,317	1,317	1,317

Table 4: **Timing and Placebo Tests around the 2005 Salience Shock**
This table reports timing and placebo tests for the continuous difference-in-differences design. Columns (1)–(3) use matched rolling three-year residual windows ending one year before the candidate break date and are estimated on the pooled sample over the 2000–2012 window. Column (1) uses Post-2002 as an information-availability benchmark. Columns (2) and (3) use placebo breaks in 2003 and 2004. Column (4) reports the actual 2005 Morningstar salience shock using the preferred pooled full-sample baseline specification, so it exactly matches the preferred all-funds specification in Table 3. All columns use fund fixed effects, style×year fixed effects, and the same lagged controls as in Table 3. Standard errors are clustered at the fund level. t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)
Tax inefficiency (1999–2001) × Post-2002	0.348			
	(0.338)			
Tax inefficiency (2000–2002) × Post-2003		0.694		
		(0.684)		
Tax inefficiency (2001–2003) × Post-2004			6.121*	
			(1.946)	
Tax inefficiency (2002–2004) × Post-2005				-5.748**
				(-1.980)
Fund FE	Yes	Yes	Yes	Yes
Style×Year FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Observations	11,576	12,722	13,491	27,663
Unique funds	1,120	1,210	1,263	1,317

Table 5: **Decomposing the Post-2005 Flow Penalty**
This table decomposes pre-period tax burden into its raw, predictable, and abnormal components in the continuous difference-in-differences framework. The dependent variable is annual net fund flow. The predictable component is the fitted value from the annual first-stage tax-burden regression, and the abnormal component is the corresponding residual. All treatment variables are averaged over 2002–2004 and expressed in percentage points. All columns use the preferred specification with fund fixed effects, style×year fixed effects, and the same lagged controls as in Table 3. Standard errors are clustered at the fund level. t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)
Pre raw tax burden (2002–2004) × Post-2005	0.750			
	(0.352)			
Pre predictable tax burden (2002–2004) × Post-2005		13.617***		17.005***
		(3.347)		(4.077)
Pre-period tax inefficiency (2002–2004) × Post-2005			-5.748**	-9.094***
			(-1.980)	(-3.028)
Fund FE	Yes	Yes	Yes	Yes
Style×Year FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Observations	27,663	27,663	27,663	27,663
Unique funds	1,317	1,317	1,317	1,317

Table 6: **Post-2005 Flow Sensitivity to Lagged Tax Burden**
This table reports a difference-in-slopes specification using lagged tax-burden measures. The dependent variable is annual net fund flow. Each column reports the pre-2005 flow sensitivity and the incremental change after 2005, captured by the interaction of the lagged tax measure with Post-2005. Columns (1) and (2) use raw tax burden measured as the lagged one-year value and the lagged three-year average, respectively. Columns (3) and (4) use abnormal tax burden measured analogously from the annual first-stage residual. All columns use the broader sample of active U.S. equity funds observed in 2004 and absorb fund fixed effects and style×year fixed effects. Controls include lagged return, lagged return squared, lagged log fund size, lagged log fund age, lagged expense ratio, lagged turnover, lagged volatility, and lagged flow. Standard errors are clustered at the fund level. t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)
Raw TB (1y)	1.135***			
	(3.323)			
Raw TB (1y) × Post-2005	-2.165***			
	(-5.240)			
Raw TB (3y)		1.730***		
		(3.017)		
Raw TB (3y) × Post-2005		-3.979***		
		(-5.707)		
Abnormal TB (1y)			1.009***	
			(2.677)	
Abnormal TB (1y) × Post-2005			-1.849***	
			(-4.246)	
Abnormal TB (3y)				2.186***
				(3.103)
Abnormal TB (3y) × Post-2005				-5.428***
				(-6.231)
Return (%)	1.436***	1.430***	1.448***	1.436***
	(20.106)	(20.138)	(20.438)	(20.339)
Return squared	-0.007***	-0.007***	-0.007***	-0.007***
	(-7.189)	(-7.031)	(-7.203)	(-6.991)
Log fund size	-17.765***	-17.929***	-17.875***	-18.021***
	(-23.991)	(-23.937)	(-24.122)	(-24.102)
Log fund age	-7.163***	-7.052***	-7.323***	-6.912***
	(-5.943)	(-5.965)	(-5.991)	(-5.853)
Expense ratio (%)	-8.111***	-7.771***	-7.512***	-7.439***
	(-3.144)	(-3.016)	(-2.879)	(-2.859)
Turnover	0.989	0.976	1.350	1.413
	(1.066)	(1.049)	(1.439)	(1.520)
Volatility (%)	0.359**	0.367**	0.358**	0.357**
	(2.446)	(2.544)	(2.447)	(2.475)
Flow (%)	0.206***	0.205***	0.200***	0.199***
	(15.159)	(15.252)	(15.935)	(16.055)
Fund FE	Yes	Yes	Yes	Yes
Style×Year FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Observed in 2004	Yes	Yes	Yes	Yes
Unique funds	1,499	1,500	1,499	1,500
Observations	30,032	30,112	29,651	29,730

Table 7: **Treatment Effects Heterogeneity by Fund Distribution Channel: Subsample Analysis**

This table reports split-sample estimates of the continuous difference-in-differences specification for direct and indirect funds. Direct funds have no positive sales load and average 12b-1 fees no greater than 0.25% in the pre-shock period. Indirect funds have a positive load or average 12b-1 fees above that threshold. Both columns report the preferred specification with lagged controls, fund fixed effects, and style×year fixed effects. Standard errors are clustered at the fund level. t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)
	Direct funds	Indirect funds
Pre-period tax inefficiency × Post-2005	1.396	-12.816***
	(0.278)	(-3.777)
Fund FE	Yes	Yes
Style×Year FE	Yes	Yes
Controls	Yes	Yes
Observations	11,064	16,591
Unique funds	523	794

Table 8: **Heterogeneity by Fund Distribution Type: Triple-Interaction Estimates**

This table reports pooled estimates that compare the post-2005 flow penalty across direct and indirect funds. Column (1) repeats the preferred baseline specification. Column (2) adds the triple interaction between pre-period tax inefficiency, Post-2005, and an indicator for indirect funds. In Column (2), the coefficient on pre-period tax inefficiency × Post-2005 is the direct-fund effect, and the triple interaction is the incremental indirect-fund effect relative to direct funds. All regressions use the preferred lagged controls, fund fixed effects, and style×year fixed effects. Standard errors are clustered at the fund level. t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)
	Baseline	Triple interaction
Pre-period tax inefficiency × Post-2005	-5.748**	0.719
	(-1.980)	(0.141)
Pre-period tax inefficiency × Post-2005 × Indicator(Indirect fund)		-12.060**
		(-2.043)
Fund FE	Yes	Yes
Style×Year FE	Yes	Yes
Controls	Yes	Yes
Observations	27,663	27,663
Unique funds	1,317	1,317

Table 9: **Subperiod Robustness of the Post-2005 Flow Penalty**
This table reports the preferred continuous difference-in-differences specification for the pooled sample and for the direct and indirect subsamples, estimated both on the full sample and on the narrower 2000–2012 window. All columns use the same preferred lagged controls, fund fixed effects, and style×year fixed effects as in the main tables. The subperiod estimates focus identification on the years closest to the 2005 disclosure change. Standard errors are clustered at the fund level. t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	Pooled Full sample	Pooled 2000–2012	Direct Full sample	Direct 2000–2012	Indirect Full sample	Indirect 2000–2012
Pre-period tax inefficiency × Post-2005	-5.748**	-9.955***	1.396	-1.922	-12.816***	-19.298***
	(-1.980)	(-3.337)	(0.278)	(-0.440)	(-3.777)	(-5.263)
Fund FE	Yes	Yes	Yes	Yes	Yes	Yes
Style×Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	27,663	14,214	11,064	5,716	16,591	8,498
Unique funds	1,317	1,315	523	521	794	794

Table 10: **Broker Classification Robustness**
This table re-estimates the indirect-fund specification on the 2000–2012 window under alternative definitions of broker or indirect distribution. Column (1) is the main definition used in the paper: any sales load in 2004 or a 12b-1 fee above 25 basis points. The remaining columns vary this rule using front-end loads only, any load in 2004, the 12b-1 threshold alone, a high-load threshold above 2%, and an average 2002–2004 rule. All columns use the preferred lagged controls, fund fixed effects, and style×year fixed effects. Standard errors are clustered at the fund level. t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	Load or 12b-1>25bp	Front-load only	Any load in 2004	12b-1 only	Load >2%	2002–2004 avg rule
Pre-period tax inefficiency × Post-2005	-19.298***	-14.032***	-19.107***	-16.222***	-13.478**	-19.525***
	(-5.263)	(-3.178)	(-5.115)	(-2.927)	(-2.394)	(-5.290)
Fund FE	Yes	Yes	Yes	Yes	Yes	Yes
Style×Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	8,498	4,869	8,196	5,215	5,300	8,671
Unique funds	794	458	767	489	506	810

Table 11: **Sample Construction Diagnostics**
This table summarizes the main sample construction steps used in the tax-salience analysis. The baseline continuous difference-in-differences sample is more restrictive because it requires a fund-level treatment based on usable residual tax burden measures in the full 2002–2004 pre-period. The broader sample used in Table 6 requires only that a fund be observed in 2004. The counts in this table describe sample construction at key filtering stages. Final regression samples can be smaller because the estimation tables additionally require non-missing lagged controls and full regression variables.

Step	Number of funds
All funds in dataset	3,321
Funds with any observation in 2004	1,739
Funds in baseline sample in 2004	1,502
Funds with usable tax-burden variable in 2004	1,532